V.B.i – Photon & Event Databases

Tom Stephens – GSSC Programmer

Outline

- Database Descriptions
- Database Requirements
- Design Goals
- D1/D2 Database System Design
  - Queue Manager
  - Event and Photon Database Servers
  - Internal Data Storage
  - Pointing & Livetime History Database Server
- Completion Milestones
- RFA Responses
- Summary

Relevant Documents

- GSSC Functional Requirements Document (433-RQMT-002)
- GSSC Design Document (GSSC-0003)
- LAT Event Summary Database Requirements document (LESDR) (GSSC-0006)
- Standard Analysis Environment Database Requirements Document (SAEDR) (GSSC-0007)

Database Descriptions

- Photon Database (D1ph)
  - Contains all events deemed to be photons by the LAT level 1 processing pipeline (~3Hz)
  - Contains all data (energy, arrival direction & time, etc.) necessary for scientific analysis and calculating IRFs
- Event Database (D1ev)
  - Contains all event data reconstructed by the LAT level 1 processing pipeline – i.e. photons and particles (~300Hz)
  - Contains additional, detailed information about each event.
  - Expected to be rarely accessed.
- Pointing and Livetime History Database (D2)
  - Contains information on the position and attitude of the spacecraft
  - Contains information about mode, status and livetime of the LAT instrument

D1 Search Definitions

- "Standard" Search
  - 15° radius circle or 30° x 30° box on the sky for a time period of one year (LESDR 5.2.3.1.1)
  - For photon database this is 200-600 MBytes of data depending on sky position

- "Large" Search
  - Photon database: Search that would return more than 2GBytes of data per year of observation (LESDR 5.2.3.1.4)
  - Event database: Search that would return more than 20Gbytes of data (LESDR 5.2.4.1.4)

D1 Database Design Requirements

- Search Parameters
  - Search on values that are real or integer numbers, Booleans, dates and times. (LESDR 5.2.1.1.1)
  - Times searchable to microsecond precision (LESDR 5.2.1.1.2)
  - 2-D positions on sphere (LESDR 5.2.1.1.3)
  - Data quality (LESDR 5.2.1.1.4)
- The database must be remotely accessible. (LESDR 5.2.1.3)
- Portability – must not be tied to a single architecture or software system. (LESDR 5.2.1.6)
- HEASARC Compatibility
  - Database will be turned over to HEASARC at the end of mission (LESDR 5.2.2.1)
  - Must not require excessive effort (>1 FTE) to maintain. (LESDR 5.2.2.1.1)

D1 Photon Database Performance Requirements

- Derived from statistics of current satellite data archives
- Search Speeds
  - Standard Search – Data returned within 30 minutes per year of data searched. (LESDR 5.2.3.1.2)
  - Standard Search with additional sub-selections – All data returned within 45 minutes per year of data searched. (LESDR 5.2.3.1.3)
  - Large Search – All data returned within 3 days. Allows for processing during off peak hours. (LESDR 5.2.3.1.4)
- Number of Requests
  - Must perform up to 60 standard searches a day. (LESDR 5.2.3.2.1)
- Data Ingest
  - Ingest of new data must be complete within 10 minutes for a 5 hour observation data set (LESDR 5.2.3.3)
  - Ingest of reprocessed data may interrupt database access for no more than 60 minutes for a 5 hour observation data set. (LESDR 5.2.3.4)
- Database Restoration
  - Must be able to restore database after a crash in <3 days per year of data (LESDR 5.2.3.4.3)

D1 Event Database Performance Requirements

- Search Speeds
  - Standard Search – All data returned within 10 hours per year of data searched. (LESDR 5.2.4.1.2)
  - Standard Search with additional sub-selections – All data returned within 15 hours per year of data searched. (LESDR 5.2.4.1.3)
  - Large Search – All data returned within 7 days. (LESDR 5.2.4.1.4)
- Number of Requests
  - Must be able to perform up to 1 standard search a day. (LESDR 5.2.4.2)
- Data Ingest
  - Ingest of new data must be complete within 100 minutes for a 5 hour observation data set. (LESDR 5.2.4.3)
  - Ingest of reprocessed data may interrupt database access for no more than 10 hours for a 5 hour observation data set. (LESDR 5.2.4.4)
- Database Restoration
  - Must be able to restore database after a crash in <1 week per year of data (LESDR 5.2.4.4.3)
- Requirements are generous and design goals provide better performance

D2 Database Design Requirements

- Search Speed
  - Retrieve 6 months of consecutive data (~50 MBytes) in 1 minute (SAEDR5.4.1.5.2)
- Number of searches
  - Must be able to handle >1500 searches a day (SAEDR 5.4.1.5.2.3)
- Data Ingest
  - Ingest of new data (5 hours of spacecraft operation) in 1 minute (SAEDR 5.4.1.5.1)
  - Ingest of reprocessed data (5 hour period) in 5 minutes (SAEDR 5.4.1.5.3)
- Database Restoration
  - Must be able to restore database after a crash in <1 day (SAEDR 5.4.1.5.4)

Design Goals

	Design Requirement	Design Goals	Current Performance
Standard D1 photon search – 1 year of data	30 min	1 min	~40 sec
Standard D1 event search – 1 year of data	10 hrs	30 min	N/T
D2 search – 6 months of data	60 sec	60 sec	7 sec
D1 photon ingest, new data – 5 hours of data	10 min	2 min	N/T
D1 event ingest, new data – 5 hours of data	100 min	20 min	N/T
D1 photon ingest, reprocessed data – 5 hours of data	60 min	12 min	N/T
D1 event ingest, reprocessed data – 5 hours of data	10 hrs	2 hrs	N/T
D2 Ingest, new data – 5 hours of data	1 min	1 min	10 sec
D2 Ingest, reprocessed data – 5 hours of data	5 min	1 min	N/T
D1 photon Data Restoration – year of data	3 days	3 hrs	10 min
D1 event Data Restoration – year of data	7 days	3 days	N/T
D2 Data Restoration – entire database	1 day	1 hr	10 min

Database Design Considerations

- Databases are large
  - D1 photons – ~300,000/day = ~110 million photons a year
  - D1 events – ~100x the number of photons + 2-10x amount of data per event
  - D2 – 1 entry every 30 seconds = ~1 million entries a year
- Did trade studies to find the fastest search methods:
  - Winner was row-filtering FITS files with CFITSIO.
  - Databases search event data stored as tables in FITS format
  - Search speed is directly related to amount of data per event
- Typical search will be a region of the sky
  - dividing the data into smaller regions would reduce the amount of data to be processed for any given search

D1/D2 Database System Design

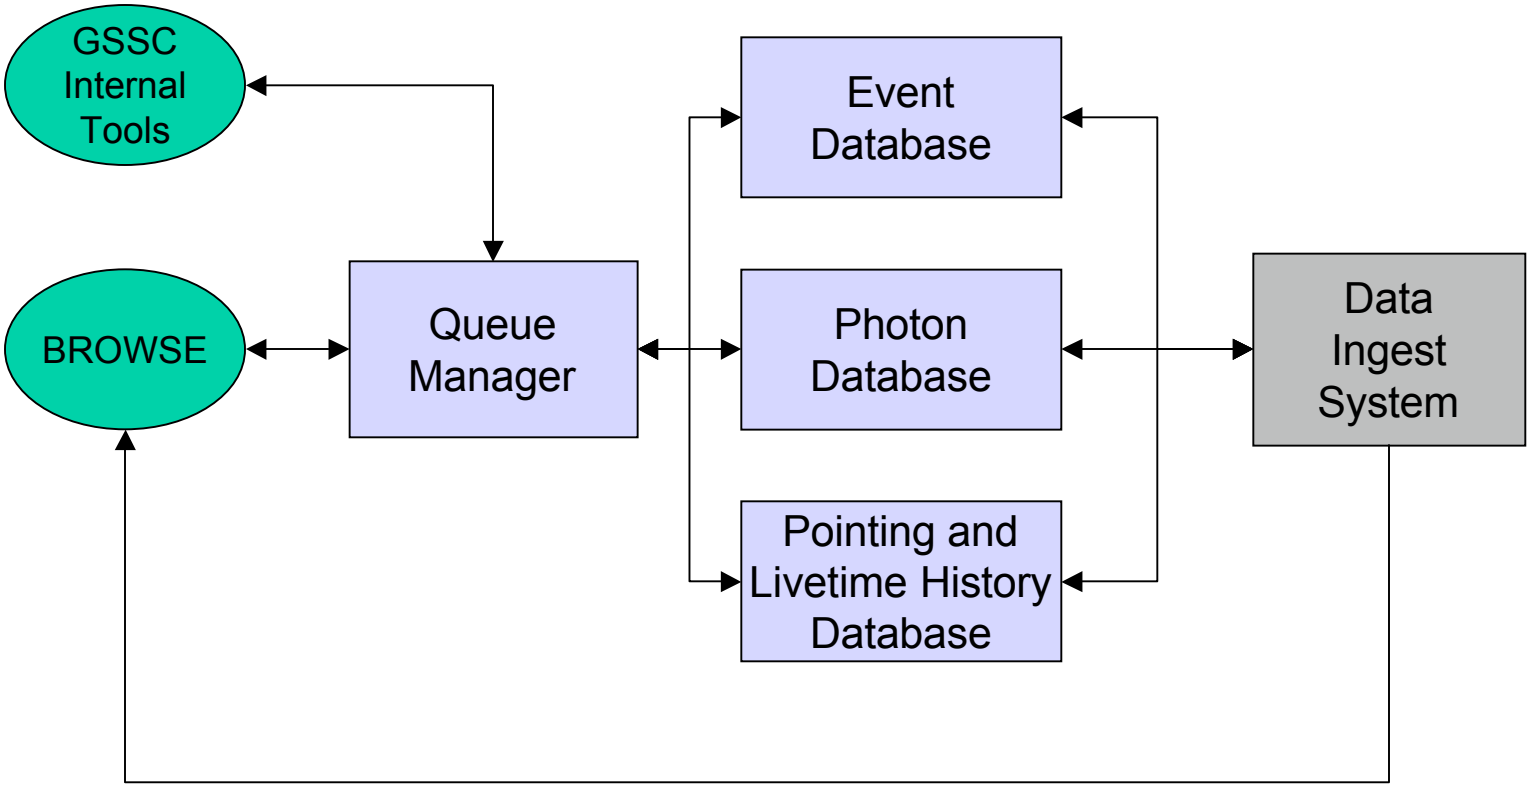

System Components

- BROWSE – User interface – HEASARC's data query system
- GSSC Internal Tools – Tools for internal manipulation of the database for testing, development and database maintenance.
- Queue Manager – handles queries from BROWSE
- Event Database – Server for the full event data
- Photon Database – Server for the photon data
- Pointing and Livetime History Database – Server for spacecraft data
- Ingest Process – feeds data to the database servers and updates BROWSE metadata

Queue Manager

- Receives all query requests from the BROWSE web interface and command line query tool
- Assigns priority based on size of query
- Sends query requests to the appropriate database servers
- Collects results from individual databases
- Returns compiled results to user
- Notifies user of estimated time to completion of query
- Notifies user of system errors preventing or delaying query completion – i.e. system down, ingest in progress, etc.

D1 Photon Database Design

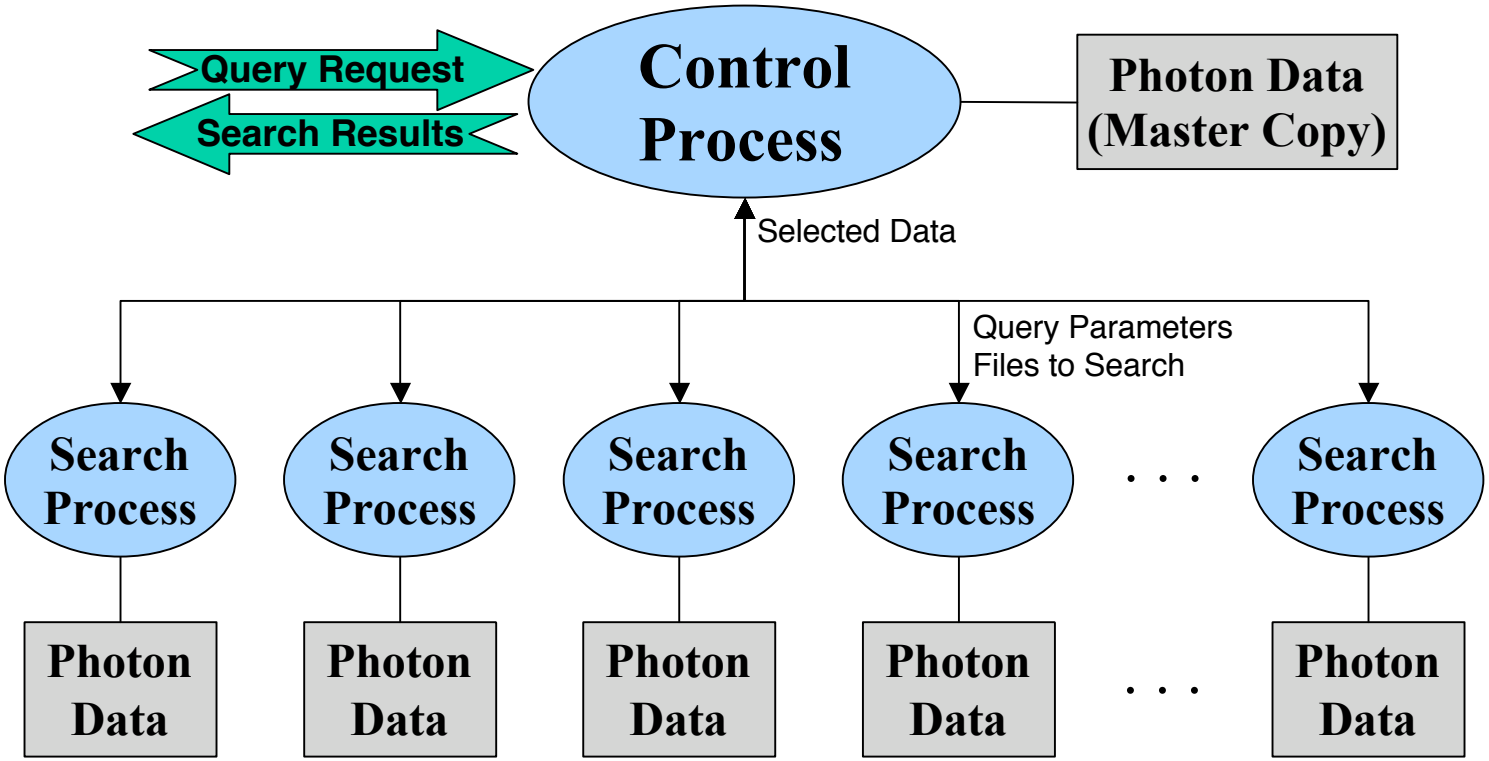

Search Engine Description

- Control process parses query requests extracting relevant keywords.
- Control process builds list of internal files to search.
- File list is distributed as needed among search processes.
  - Number of search nodes to use depends on size of search and current utilization
  - Control process tries to optimize search speed/files returned
- Search performed by each process on its assigned files
  - Results from each processed stored in single merged file
  - Filename returned to control process
- Final results compiled by control process and sent to Queue Manager.

Photon Database Internal Storage

- All data is in HEASARC compatible FITS files
- Each node (control and search) has a complete copy of the photon data.
  - Fast data access from internal disk
  - Multiple backups in case of failure of a single data disk
- Data broken into sky regions and time periods in internal data files
- Hierarchical Triangular Mesh (HTM) used to define regions
  - Developed for Sloan Digital Sky Survey at Johns Hopkins
  - Recursively divides sky into spherical triangles
- Conducted trade study to determine optimal combination of HTM pixelization level and time binning
  - Best time of ~39 sec was level 3 pixelization (512 sky regions) with 2 month time bins

Internal Storage Trade Study

- Explored a grid of HTM pixelization level and time bins
  - HTM Levels 2, 3, 4, 5 & 6
  - Time bins ranging from 1 day to 1 year
- Performed battery of searches to determine average search time.

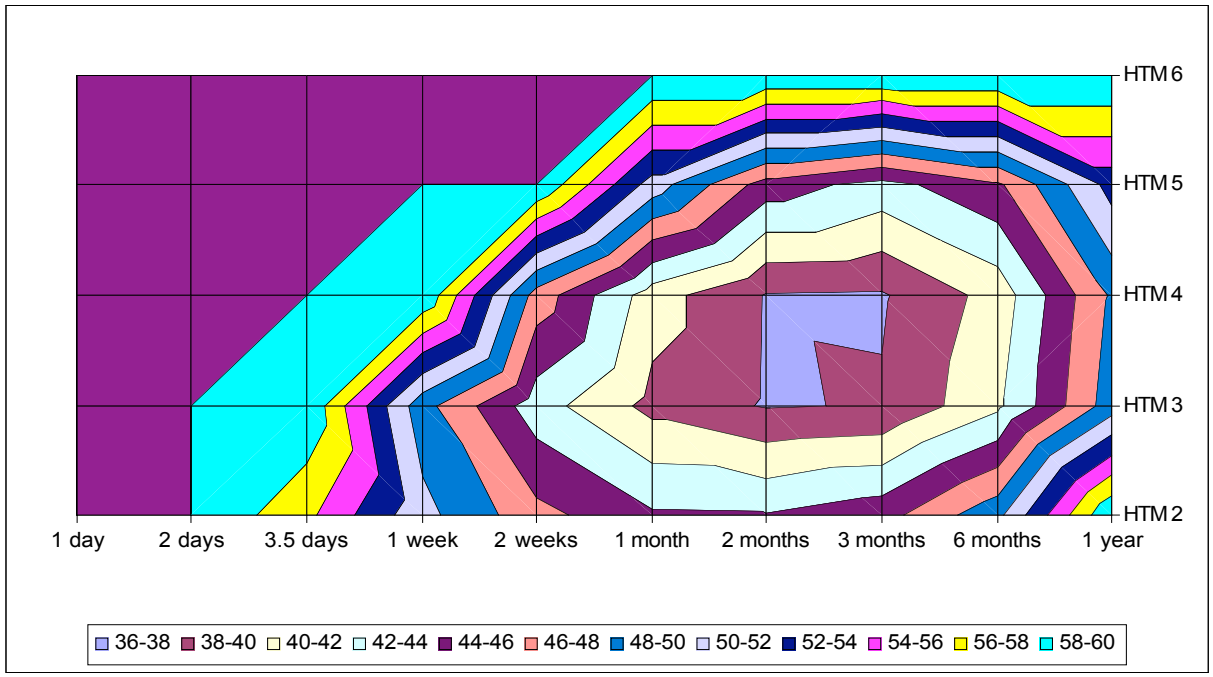

D2 Database Design

- Basic functionality identical to Event and Photon database

- Processing not nearly as difficult – only a single process
  - Only has to search on time – 1D instead of spatial (2D) search
  - Much smaller data volume – ~100x smaller than Photon database
  - Smaller number of data columns – ~1/2 the number in Photon database

Milestones

- Prototypes tested as part of the LAT Data Challenge I (Dec '03 – Feb '04) – Successfully served one week of data to the LAT team throughout the Data Challenge period.

- Preliminary production version to be used as part of LAT Data Challenge II (Spring '05)
  - Queue Manager and Ingest Modules ready by 08/06/04
  - D1 Photon Server ready by 09/17/04
  - D2 Server ready by 10/29/04
  - Web Interface ready by 12/17/04
  - D1 Event Server will not be used

- Full system to be completed by 6/15/05 in preparation for GRT 5 (11/15/05)

RFA #14

- RFA #14 – The size of the staging area for users of the GLAST archive needs to be defined, and a contingency plan needed to deal with problems if the stage disk fills up. The impact of filling the stage disk on access to GLAST data needs to be addressed.

- Size of staging area
  - Using current definitions of file size and data persistence time we estimate a need for ~100 GB of staging space for the photon database
  - Current system configuration includes 1TB of staging space
- Filling of the staging disk
  - Temporarily shortening the persistence time allows older data to be removed sooner
  - Database system will monitor available resources

RFA #15

- RFA #15 – Searches of the GLAST archive need to be reproducible; this is a problem for GLAST since the photons stored in the archive are not static but may change with reprocessing. This is also a problem since derived calibration products like IRFs will depend on photon properties and hence on processing software.

- Calibration Issue
  - All photons have a calibration software version number associated with them to link them to the proper IRFs during data analysis
  - All versions of the IRFs are part of the CALDB
- Reprocessing Issue
  - When reprocessed data is ingested, the state of the database will be saved and tagged with a version number and date.
  - Old files moved to archival area
  - If old data is needed, the state of the database can be restored and the query performed

RFA #17

- RFA #17 – The design of the D1 database includes no provision for proprietary protection of data. This is entirely consistent with current mission policy, but there's a chance that this could be revised to be consistent with the practice for every other observatory-class mission.

- "We are not aware of any parties pushing for a real proprietary data policy, and therefore detailed planning to support such a policy is not warranted." However,
- Data to be restricted could be marked with a flag to indicate that it is "proprietary".
- Restricting access to data can be then achieved by requiring a password to access data that is marked as "proprietary" in the database .
- As the data becomes unrestricted the flag is changed to indicate a "public" status.

Summary

- Requirements have driven design of primary databases
- Data stored in HEASARC-compatible FITS files
- Prototypes met requirements and design goals
- Development of production versions underway
- Many specific milestones and functional tests during development leading up to official testing in latter half of 2005.